

Clustering RDF Databases Using Tunable-LSH

Güneş Aluç, M. Tamer Özsu, and Khuzaima Daudjee

University of Waterloo David R. Cheriton School of Computer Science
 {galuc, tamer.ozsu, kdaudjee}@uwaterloo.ca

ABSTRACT

The Resource Description Framework (RDF) is a W3C standard for representing graph-structured data, and SPARQL is the standard query language for RDF. Recent advances in Information Extraction, Linked Data Management and the Semantic Web have led to a rapid increase in both the volume and the variety of RDF data that are publicly available. As businesses start to capitalize on RDF data, RDF data management systems are being exposed to workloads that are far more diverse and dynamic than what they were designed to handle. Consequently, there is a growing need for developing workload-adaptive and self-tuning RDF data management systems. To realize this vision, we introduce a fast and efficient method for dynamically clustering records in an RDF data management system. Specifically, we assume nothing about the workload upfront, but as SPARQL queries are executed, we keep track of records that are co-accessed by the queries in the workload and physically cluster them. To decide dynamically (hence, in constant-time) where a record needs to be placed in the storage system, we develop a new locality-sensitive hashing (LSH) scheme, TUNABLE-LSH. Using TUNABLE-LSH, records that are co-accessed across similar sets of queries can be hashed to the same or nearby physical pages in the storage system. What sets TUNABLE-LSH apart from existing LSH schemes is that it can auto-tune to achieve the aforementioned clustering objective with high accuracy even when the workloads change. Experimental evaluation of TUNABLE-LSH in our prototype RDF data management system, *chameleon-db*, as well as in a standalone hashtable shows significant end-to-end improvements over existing solutions.

1. INTRODUCTION

Physical data organization plays an important role in the performance tuning of database management systems. A particularly important problem is clustering (in the storage system) records that are frequently co-accessed by queries in a workload. Suboptimal clustering has negative performance implications due to random I/O and cache stalls [5]. This problem has received attention in the context of SQL databases and has led to the introduction of tuning

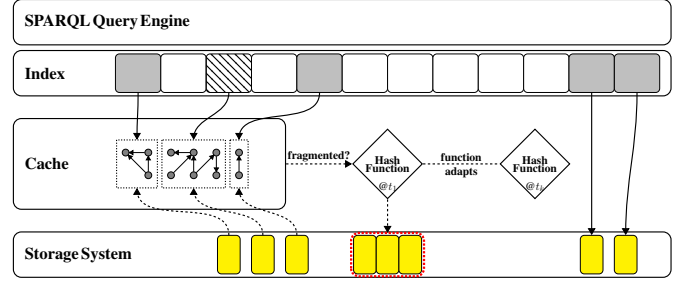


Figure 1: Adaptive record placement using a combination of adaptive hashing and caching [9].

advisors that work either in an *offline* [4, 59] or *online* fashion (i.e., self-tuning databases) [24].

In this paper, we address the problem in the context of RDF data management systems. SPARQL workloads are far more dynamic than SQL workloads [13, 43]; yet, tuning techniques for RDF data management systems are in their infancy, and relational solutions are not directly applicable. More specifically, depending on the workload, it might be necessary to completely change the underlying physical representation in an RDF data management system, such as by dynamically switching from a row-oriented representation to a columnar representation [9]. On the other hand, existing online tuning techniques work well only when the schema changes are minor [20]. Consequently, with the increasing demand to support highly dynamic workloads in RDF [13, 43], there is a growing need to develop more adaptive tuning solutions, in which records in an RDF database can be dynamically and continuously clustered based on the current workload.

Whenever a SPARQL query is executed, there is an opportunity to observe how records in an RDF database are being utilized. This information about query access patterns can be used to dynamically cluster records in the storage system. Dynamism is important in RDF systems because of the high variability and dynamism in SPARQL workloads [13, 43]. While this problem has been studied as physical clustering [47] and distribution design [22], the highly dynamic nature of the queries over RDF data introduces new challenges. First, traditional algorithms are offline, and since clustering is an NP-hard problem and most approximations have quadratic complexity [42], they are not suitable for online database clustering. Instead, techniques are needed with similar clustering objectives, but that have constant running time. Second, systems are typically expected to execute most queries in subseconds [50], leaving only fractions of a second to update their physical data structures (i.e., in our case, we are concerned with dynamically moving

records across the storage system).

We address the aforementioned issues by making two contributions. First, as shown in Fig. 1, instead of clustering the whole database, we cluster only the “warm” portions of the database by relying on the admission policy of the existing database cache. Second, we develop a self-tuning locality-sensitive hash (LSH) function, namely, TUNABLE-LSH to decide in constant-time where in the storage system to place a record. TUNABLE-LSH has two important properties:

- It tries to ensure that (i) records with *similar* utilization patterns (i.e., those records that are co-accessed across similar sets of queries) are mapped as much as possible to the same pages in the storage system, while (ii) minimizing the number of records with *dissimilar* utilization patterns that are falsely mapped to the same page.
- Unlike conventional LSH [30,40], TUNABLE-LSH can auto-tune so as to achieve the aforementioned clustering objectives with high accuracy even when the workloads change.

These ideas are illustrated in Fig. 1. Let us assume that initially, the records in a database are *not* clustered according to any particular workload. Therefore, the performance of the system is sub-optimal. However, every time records are fetched from the storage system, there is an opportunity to bring together into a single page those records that are co-accessed but are fragmented across the storage system. TUNABLE-LSH achieves these with minimal overhead. Furthermore, TUNABLE-LSH is continuously updated to reflect any changes in the workload characteristics. Consequently, as more queries are executed, records in the database become more clustered, and the performance of the system improves.

The paper is organized as follows: Section 2 discusses related work. Section 3 gives a conceptual description of the problem. Section 4 describes the overview of our approach while Section 5 provides the details. In Section 6, we describe how physical clustering takes place in the database, in particular, how TUNABLE-LSH can be used in an RDF data management system, and we experimentally evaluate our techniques. Finally, we discuss conclusions and future work in Section 7.

2. RELATED WORK

Locality-sensitive hashing (LSH) [30,40] has been used in various contexts such as nearest neighbour search [12,14,28,36,40,56], Web document clustering [18,19] and query plan caching [7]. In this paper, we use LSH in the physical design of RDF databases. While multiple families of LSH functions have been developed [19,23,25,30,40], these functions assume that the input distribution is either uniform or static. In contrast, TUNABLE-LSH can continuously adapt to changes in the input distribution to achieve higher accuracy, which translates to adapting to changes in the query access patterns in the workloads in the context of RDF databases.

Physical design has been the topic of an ongoing discussion in the world of RDF and SPARQL [1,9,55,58]. One option is to represent data in a single large table [21] and build clustered indexes, where each index implements a different sort order [34,52,57]. It has also been argued that grouping data can help improve performance [1,55]. For this reason, multiple physical representations have been developed: in the *group-by-predicate* representation, the database is vertically partitioned and the tables are stored in a column-store [1]; in the *group-by-entity* representation, implicit relationships within the database are discovered (either manually [58] or automatically [17]), and the RDF data are mapped to a relational database; and in the *group-by-vertex* representation,

RDF’s inherent graph-structure is preserved, whereby data can be grouped on the vertices in the graph [60]. These workload-oblivious representations have issues for different types of queries, due to reasons such as fragmented data, unnecessarily large intermediate result tuples generated during query evaluation and/or suboptimal pruning by the indexes [9].

To address some of these issues, workload-aware techniques have been proposed [31,35]. For example, view materialization techniques have been implemented for RDF over relational engines [31]. However, these materialized views are difficult to adapt to changing workloads for reasons discussed in Section 1. Workload-aware distribution techniques have also been developed for RDF [35] and implemented in systems such as WARP [35] and Partout [29], but, these systems are not runtime-adaptive. With TUNABLE-LSH, we aim to address the problem adaptively, by clustering fragmented records in the database based on the workload.

While there are self-tuning SQL databases [24,38,39] and techniques for automatic schema design in SQL [4,15,47,59], these techniques are not directly applicable to RDF. In RDF, the advised changes to the underlying physical schema can be drastic, for example, requiring the system to switch from a row-oriented representation to a columnar one, all at runtime, which are hard to achieve using existing techniques. Consequently, there have been efforts in designing workload-adaptive and self-tuning RDF data management systems [6,9,10,53,54]. In H2RDF [53], the choice between centralized versus distributed execution is made adaptively. In PHDStore [6], data are adaptively replicated and distributed across the compute nodes; however, the underlying physical layout is fixed within each node. A mechanism for adaptively caching partial results is introduced in [54]. With TUNABLE-LSH, we are trying to address the adaptive record layout problem, therefore, we believe that TUNABLE-LSH will complement existing techniques and facilitate the development of runtime adaptive RDF systems.

3. PRELIMINARIES

Given a sequence of database records that represent the records’ serialization order in the storage system, the access patterns of a query can conceptually be represented as a bit vector, where a bit is set to 1 if the corresponding record in the sequence is accessed by the query. We call this bit vector a *query access vector* (\vec{q}).

Depending on the system, a record may denote a single RDF triple (i.e., the atomic unit of information in RDF), as in systems like RDF-3x [51], or a collection of RDF triples such as in chameleon-db [10,11]. Our conceptual model is applicable either way.

As more queries are executed, their query access vectors can be accumulated column-by-column in a matrix, as shown in Fig. 2a. We call this matrix a *query access matrix*. For presentation, let us assume that queries are numbered according to their order of execution by the RDF data management system.

Each row of the query access matrix constitutes what we call a *record utilization vector* (\vec{r}), which represents the set of queries that access record r . As a convention, to distinguish between a query and its access vector (likewise, a record and its utilization vector), we use the symbols q and \vec{q} (likewise, r and \vec{r}), respectively. The complete list of symbols are given in Table 1.

To model the memory hierarchy, we use an additional notation in the matrix representation: records that are physically stored together on the same disk/memory page should be grouped together in the query access matrix. For example, Fig. 2a and Fig. 2b represent two alternative ways in which the records in an RDF database can be clustered (groups are separated by horizontal dashed lines). Even though both figures depict essentially the same query access

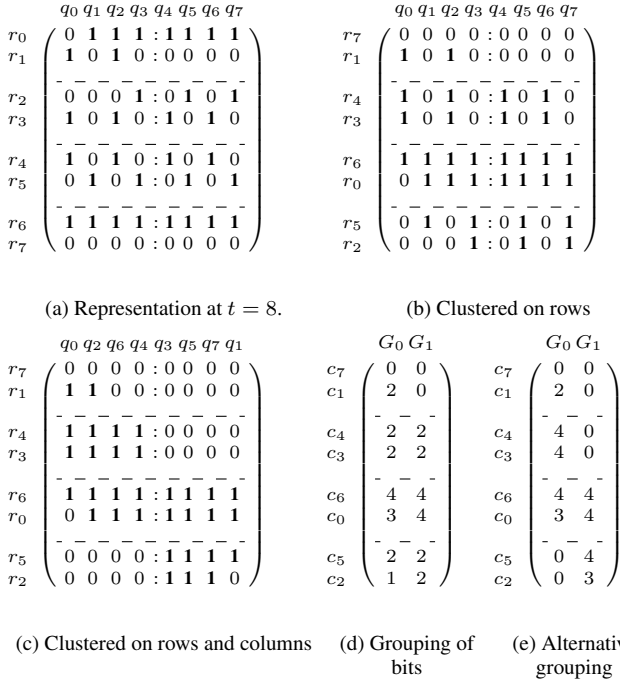


Figure 2: Matrix representation of query access patterns.

patterns, the physical organization in Fig. 2b is preferable, because in Fig. 2a, most queries require access to 4 pages each, whereas in Fig. 2b, the number of accesses is reduced by almost half.

Given a sequence of queries and the number of pages in the storage system, our objective is to *store records with similar utilization vectors together so as to minimize the total number of page accesses*. To determine the similarity between record utilization vectors, we rely on the following property. Two records are co-accessed by a query if both of the corresponding bits in that query's access vector are set to 1. Extending this concept to a set of queries, we say that two records are co-accessed across multiple queries if the corresponding bits in the record utilization vectors are set to 1 for all the queries in the set. For example, according to Fig. 2a, records r_1 and r_3 are co-accessed by queries q_0 and q_2 , and records r_0 and r_6 are co-accessed across the queries q_1 – q_7 .

Given a sequence of queries, it may often be the case that a pair of records are not co-accessed in *all* of the queries. Therefore, to measure the extent to which a pair of records are co-accessed, we rely on their Hamming distance [32]. Specifically, given two record utilization vectors for the same sequence of queries, their Hamming distance—denoted as $\delta(\vec{q}_x, \vec{q}_y)$ —is defined as the minimum number of substitutions necessary to make the two bit vectors the same [32].¹ Hence, the smaller the Hamming distance between a pair of records, the greater the extent to which they are co-accessed.

Consider the record utilization vectors \vec{r}_0 , \vec{r}_2 , \vec{r}_5 and \vec{r}_6 across the query sequence q_0 – q_7 in Fig. 2a. The pairwise Hamming distances are as follows: $\delta(r_0, r_6) = 1$, $\delta(r_2, r_5) = 1$, $\delta(r_0, r_5) = 3$, $\delta(r_0, r_2) = 4$, $\delta(r_5, r_6) = 4$ and $\delta(r_2, r_6) = 5$. Consequently, to achieve better physical clustering, we should try to store r_0 and r_6 together and r_2 and r_5 together, while keeping r_0 and r_6 apart from r_2 and r_5 .

¹The Hamming distance between two record utilization vectors is equal to their edit distance [46], as well as the Manhattan distance [44] between these two vectors in l_1 norm.

	Symbol	Description
Constants	ω	database size (i.e., number of records)
	ϵ	number of pages in the storage system
	k	maximum no. of query access vectors that can be stored
	b	number of entries in each record utilization counter
	t	current time
Data structures	\vec{q}	query access vector (contains ω bits)
	\vec{r}	record utilization vector (contains k bits)
	\vec{c}	record utilization counter (contains b entries)
	\vec{P}	depending on the context, a point in a k -dimensional or b -dimensional (Taxicab) space
	$M_{\omega \times k}$	query access matrix; contains the last k most representative query access vectors (in columns), or equivalently, ω record utilization vectors (in rows)
	$C_{\omega \times b}$	frequency matrix; represents record utilization frequency over b groups of query access vectors
Accessors	$q[i]$	value of the i^{th} bit in query access vector \vec{q}
	$r[i]$	value of the i^{th} bit in record utilization vector \vec{r}
	$c[i]$	value of the i^{th} entry in record utilization counter \vec{c}
	$P[i]$	value of the i^{th} coordinate in point \vec{P}
	$M[i][j]$	value of the i^{th} row and j^{th} column in matrix
Distances	$C[i][j]$	value of the i^{th} row and j^{th} column in matrix
	$\delta(\vec{r}_x, \vec{r}_y)$	Hamming distance between two record utilization vectors
	$\delta^H(\vec{q}_x, \vec{q}_y)$	MIN-HASH distance between two query access vectors
	$\delta^M(\vec{P}_x, \vec{P}_y)$	Manhattan distance between two points

Table 1: Symbols used throughout the manuscript

4. OVERVIEW OF TUNABLE-LSH

Although we are dealing with a clustering problem, the dynamic nature of queries over RDF data necessitate a solution different than existing ones [9]. That is, while conventional clustering algorithms [42] might be perfectly applicable for the *offline* tuning of a database, in an *online* scenario, even the most efficient algorithms may be impractical unless records are clustered on-the-fly and within microseconds. Clustering is an NP-complete problem [42], and most approximations take at least quadratic time. It is not very well-understood which clustering algorithm is more suitable for which types of input distributions [2], let alone the fact that incremental versions of these algorithms are largely domain-specific [3]. Therefore, we develop TUNABLE-LSH, which is a self-tuning locality-sensitive hash (LSH) function. As records are fetched from the storage system, we keep track of records that are fragmented. Then, we use TUNABLE-LSH to decide, in constant-time, how a fragmented record needs to be clustered in the storage system (cf., Fig. 1). Furthermore, we develop methods to continuously auto-tune this LSH function to adapt to changing query access patterns that we encounter while executing a workload. This way, TUNABLE-LSH can achieve much higher clustering accuracy than conventional LSH schemes, which are static.

Let $\mathbb{Z}_{\alpha \dots \beta}$ denote the set of integers in the interval $[\alpha, \beta]$, and let $\mathbb{Z}_{\alpha \dots \beta}^n$ denote the n -fold Cartesian product:

$$\underbrace{\mathbb{Z}_{\alpha \dots \beta} \times \dots \times \mathbb{Z}_{\alpha \dots \beta}}_n.$$

Let us assume that we are given a non-injective, surjective function $f : \mathbb{Z}_{0 \dots (k-1)} \rightarrow \mathbb{Z}_{0 \dots (b-1)}$, where $b \ll k$, and for all $y \in \mathbb{Z}_{0 \dots (b-1)}$, it holds that

$$|\{x : f(x) = y\}| \leq \left\lceil \frac{k}{b} \right\rceil.$$

In other words, f is a hash function with the property that, given k input values and b possible outcomes, no more than $\lceil \frac{k}{b} \rceil$ values in the domain of the function will be hashed to the same value. Then, we define TUNABLE-LSH as $h : \mathbb{Z}_{0 \dots 1}^k \rightarrow \mathbb{Z}_{0 \dots (\epsilon-1)}$, where ϵ

represents the number of pages in the storage system. More specifically, h is defined as a composition of two functions h_1 and h_2 .

DEFINITION 1 (TUNABLE-LSH).

Let

$$\vec{r} = (r[0], \dots, r[k-1]) \in \mathbb{Z}_{0 \dots 1}^k, \text{ and}$$

$$\vec{c} = (c[0], \dots, c[b-1]) \in \mathbb{Z}_{0 \dots \lceil \frac{k}{b} \rceil}^b.$$

Then, a tunable LSH function h is defined as

$$h = h_2 \circ h_1$$

where

$$h_1 : \mathbb{Z}_{0 \dots 1}^k \rightarrow \mathbb{Z}_{0 \dots \lceil \frac{k}{b} \rceil}^b, \text{ where } h_1(\vec{r}) = \vec{c} \text{ iff}$$

$$\forall y \ c[y] = \sum_{x=0}^{k-1} \begin{cases} r[x] & : f(x) = y \\ 0 & : f(x) \neq y \end{cases}$$

$$h_2 : \mathbb{Z}_{0 \dots \lceil \frac{k}{b} \rceil}^b \rightarrow \mathbb{Z}_{0 \dots (\epsilon-1)}, \text{ where } h_2(\vec{c}) = v \text{ iff}$$

$$v = \begin{cases} \text{coordinate of } \vec{c} \text{ (rounded to the} \\ \text{nearest integer) on a space-filling} \\ \text{curve [49] of length } \epsilon \text{ that covers} \\ \mathbb{Z}_{0 \dots \lceil \frac{k}{b} \rceil}^b \end{cases}$$

According to Def. 1, h is constructed as follows:

1. Using a hash function f (which can be treated as a black box for the moment), a record utilization vector \vec{r} with k bits is divided into b disjoint segments $\vec{r}_0, \dots, \vec{r}_{b-1}$ such that $\vec{r}_0, \dots, \vec{r}_{b-1}$ contain all the bits in \vec{r} , and each $\vec{r}_i \in \{\vec{r}_0, \dots, \vec{r}_{b-1}\}$ has at most $\lceil \frac{k}{b} \rceil$ bits. Then, a record utilization counter \vec{c} with b entries is computed such that the i^{th} entry of \vec{c} (i.e., $c[i]$) contains the number of 1-bits in \vec{r}_i . Without loss of generality, a record utilization counter \vec{c} can be represented as a b -dimensional point in the coordinate system $\mathbb{Z}_{0 \dots \lceil \frac{k}{b} \rceil}^b$.
2. The final hash value is computed by ordering the points in $\mathbb{Z}_{0 \dots \lceil \frac{k}{b} \rceil}^b$ using a space-filling curve [49].

In Section 5.1, we show that TUNABLE-LSH that maps k -dimensional record utilization vectors to natural numbers in the interval $[0, \dots, \epsilon - 1]$ is locality-sensitive, with two important implications: (i) records with similar record utilization vectors (i.e., small Hamming distances) are likely going to be hashed to the same value, while (ii) records with dissimilar record utilization vectors are likely going to be separated. Therefore, the problem of clustering records in the storage system can be approximated using TUNABLE-LSH, such that clustering n records takes $O(n)$ time.

The quality of TUNABLE-LSH, that is, how well it approximates the original Hamming distances, depends on two factors: (i) the characteristics of the workload so far, which is reflected by the bit distribution in the record utilization vectors, and (ii) the choice of f . In Section 5.2, we demonstrate that f can be tuned to adapt to the changing patterns in record utilization vectors to maintain the approximation quality of TUNABLE-LSH at a steady and high level.

Algorithm 1 Initialize

Ensure:

```

1: procedure INITIALIZE()
2:   construct int  $C[\omega][2b]$  ▷ For simplicity,  $C$  is allocated statically; however, in practice, it can be allocated dynamically to reduce memory footprint.

3:   for all  $i \in (0, \dots, \omega - 1)$  do
4:     for all  $j \in (0, \dots, 2b - 1)$  do
5:        $C[i][j] \leftarrow 0$ 
6:     end for
7:   end for
8: end procedure

```

Algorithm 2 Tune

Require:

\vec{q}_t : query access vector produced at time t

Ensure:

Underlying data structures are updated and f is tuned such that the LSH function maintains a steady approximation quality

```

1: procedure TUNE( $\vec{q}_t$ )
2:   RECONFIGURE- $F(\vec{q}_t)$ 
3:   for all  $i \in \text{POSITIONAL}(\vec{q}_t)$  do
4:      $\text{loc} \leftarrow f(t)$ 
5:     if  $\text{loc} < (\text{shift} \% b)$  then
6:        $\text{loc} += b$ 
7:     end if
8:      $C[i][\text{loc}]++$  ▷ Increment record utilization counters based on the new query access pattern
▷ Reset “old” counters

9:   if  $t \% \lceil \frac{k}{b} \rceil = 0$  then
10:      $\text{shift}++$ 
11:      $C[i][(\text{shift}+b)\%2b] \leftarrow 0$ 
12:   end if
13: end for
14: end procedure

```

Algorithms 1–3 present our approach for computing the outcome of TUNABLE-LSH and for incrementally tuning the LSH function every time a query is executed. Note that we have two design considerations: (i) tuning should take constant-time, otherwise, there is no point in using a function, (ii) the memory overhead should be low because it would be desirable to maximize the allocation of memory to core database functionality. Consequently, instead of relying on record utilization vectors, the algorithm computes and incrementally maintains record utilization counters (cf., Algorithm 1) that are much easier to maintain and that have a much smaller memory footprint due to the fact that $b \ll k$. Then, whenever there is a need to compute the outcome of the LSH function for a given record, the HASH procedure is called with the *id* of the record, which in turn relies on h_2 to compute the hash value (cf., Algorithm 3).

The TUNE procedure looks at the next query access vector, and updates f (line 2), which will be discussed in more detail in Section 5.2. Then it computes positions of records that have been accessed by that query (line 3), and increments the corresponding entries in the utilization counters of those records that have been accessed (line 8). To determine which entry to increment, the algorithm relies on h_1 , hence, $f(t)$ (cf., Def. 1) and a shifting scheme. In line 11, old entries in record utilization counters are reset based on an approach that we discuss in Section 5.3. In that section we also discuss the shifting scheme.

5. DETAILS OF TUNABLE-LSH AND OPTIMIZATIONS

Algorithm 3 Hash

Require:

id: id of record whose hash is being computed

Ensure:

Hash value is returned

```

1: procedure HASH(id)
2:   return Z-VALUE(C[id])           ▷ Apply  $h_2$ 
3: end procedure

```

This section is structured as follows: Section 5.1 shows that TUNABLE-LSH has the properties of a locality-sensitive hashing scheme. Section 5.2 describes our approach for tuning f based on the most recent query access patterns, and Section 5.3 explains how old bits are removed from record utilization counters.

5.1 Properties of Tunable-LSH

In this part, we discuss the locality-sensitive properties of h_1 and h_2 , and demonstrate that $h_2 \circ h_1$ can be used for clustering the records. First, we show the relationship between record utilization vectors and the record utilization counters that are obtained by applying h_1 .

THEOREM 1 (DISTANCE BOUNDS). *Given a pair of record utilization vectors \vec{r}_1 and \vec{r}_2 with size k , let \vec{c}_1 and \vec{c}_2 denote two record utilization counters with size b such that $\vec{c}_1 = h_1(\vec{r}_1)$ and $\vec{c}_2 = h_1(\vec{r}_2)$ (cf., Def. 1). Furthermore, let $c_1[i]$ and $c_2[i]$ denote the i^{th} entry in \vec{c}_1 and \vec{c}_2 , respectively. Then,*

$$\delta(\vec{r}_1, \vec{r}_2) \geq \sum_{i=0}^{b-1} |c_1[i] - c_2[i]| \quad (1)$$

where $\delta(\vec{r}_1, \vec{r}_2)$ represents the Hamming distance between \vec{r}_1 and \vec{r}_2 .

PROOF SKETCH 1. We prove Thm. 1 by induction on b .

Base case: Thm. 1 holds when $b = 1$. According to Def. 1, when $b = 1$, $c_1[0]$ and $c_2[0]$ correspond to the total number of 1-bits in \vec{r}_1 and \vec{r}_2 , respectively. Note that the Hamming distance between \vec{r}_1 and \vec{r}_2 will be smallest if and only if these two record utilization vectors are aligned on as many 1-bits as possible. In that case, they will differ in only $|c_1[0] - c_2[0]|$ bits, which corresponds to their Hamming distance. Consequently, Eqn. 1 holds for $b = 1$.

Inductive step: We show that if Eqn. 1 holds for $b \leq \alpha$, where α is a natural number greater than or equal to 1, then it must also hold for $b = \alpha + 1$. Let $\Pi_f(\vec{r}, g)$ denote a record utilization vector $r' = (r'[0], \dots, r'[k-1])$ such that for all $i \in \{0, \dots, k-1\}$, $r'[i] = r[i]$ holds if $f(i) = g$, and $r'[i] = 0$ otherwise. Then,

$$\delta(\vec{r}_1, \vec{r}_2) = \sum_{g=0}^{b-1} \delta(\Pi_f(\vec{r}_1, g), \Pi_f(\vec{r}_2, g)). \quad (2)$$

That is, the Hamming distance between any two record utilization vectors is the summation of their individual Hamming distances within each group of bits that share the same hash value with respect to f . This property holds because f is a (total) function, and Π_f masks all the irrelevant bits. As an abbreviation, let $\delta_g = \delta(\Pi_f(\vec{r}_1, g), \Pi_f(\vec{r}_2, g))$. Then, due to the same reasoning as in the base case, for $g = \alpha$, the following equation holds:

$$\delta_\alpha(\vec{r}_1, \vec{r}_2) \geq |c_1[\alpha] - c_2[\alpha]| \quad (3)$$

Consequently, using the additive property in Eqn. 2, it can be shown that Eqn 1 holds also for $b = \alpha + 1$. Thus, by induction, Thm. 1 holds. ■

Thm. 1 suggests that the Hamming distance between any two record utilization vectors \vec{r}_1 and \vec{r}_2 can be approximated using record utilization counters $\vec{c}_1 = h_1(\vec{r}_1)$ and $\vec{c}_2 = h_1(\vec{r}_2)$ because Eqn. 1 provides a lower bound on $\delta(\vec{r}_1, \vec{r}_2)$. In fact, the right-hand side of Eqn. 1 is equal to the Manhattan distance [44] between \vec{c}_1 and \vec{c}_2 in $\mathbb{Z}_{0 \dots \lceil \frac{k}{b} \rceil}^b$, and since $\delta(\vec{r}_1, \vec{r}_2)$ is equal to the Manhattan distance between \vec{r}_1 and \vec{r}_2 in $\mathbb{Z}_{0 \dots 1}^k$, it is easy to see that h_1 is a transformation that approximates Manhattan distances. The following corollary captures this property.

COROLLARY 2 (DISTANCE APPROXIMATION). *Given a pair of record utilization vectors \vec{r}_1 and \vec{r}_2 with size k , let \vec{c}_1 and \vec{c}_2 denote two points in the coordinate system $\mathbb{Z}_{0 \dots \lceil \frac{k}{b} \rceil}^b$ such that $\vec{c}_1 = h_1(\vec{r}_1)$ and $\vec{c}_2 = h_1(\vec{r}_2)$ (cf., Def. 1). Let $\delta^M(\vec{r}_1, \vec{r}_2)$ denote the Manhattan distance between \vec{r}_1 and \vec{r}_2 , and let $\delta^M(\vec{c}_1, \vec{c}_2)$ denote the Manhattan distance between \vec{c}_1 and \vec{c}_2 . Then, the following holds:*

$$\delta(\vec{r}_1, \vec{r}_2) = \delta^M(\vec{r}_1, \vec{r}_2) \geq \delta^M(\vec{c}_1, \vec{c}_2) \quad (4)$$

PROOF SKETCH 2. Hamming distance in $\mathbb{Z}_{0 \dots 1}^k$ is a special case of Manhattan distance. Furthermore, by definition [44], the right hand side of Eqn. 1 equals the Manhattan distance $\delta^M(\vec{c}_1, \vec{c}_2)$; therefore, Eqn. 4 holds. ■

Next, we demonstrate that h_1 is a locality-sensitive transformation [30, 40]. In particular, we use the definition of locality-sensitiveness by Tao et al. [56], and show that the probability that two record utilization vectors \vec{r}_1 and \vec{r}_2 are transformed into “nearby” record utilization counters \vec{c}_1 and \vec{c}_2 increases as the (Manhattan) distance between r_1 and r_2 decreases.

THEOREM 3 (GOOD APPROXIMATION). *Given a pair of record utilization vectors \vec{r}_1 and \vec{r}_2 with size k , let \vec{c}_1 and \vec{c}_2 denote two points in the coordinate system $\mathbb{Z}_{0 \dots \lceil \frac{k}{b} \rceil}^b$ such that $\vec{c}_1 = h_1(\vec{r}_1)$, $\vec{c}_2 = h_1(\vec{r}_2)$ and $b = 1$ (cf., Def. 1). Let $\delta^M(\vec{r}_1, \vec{r}_2)$ denote the Manhattan distance between \vec{r}_1 and \vec{r}_2 , and let $\delta^M(\vec{c}_1, \vec{c}_2)$ denote the Manhattan distance between \vec{c}_1 and \vec{c}_2 . Furthermore, let $\text{PR}_{\delta^M \leq \Theta}(x)$ be a shorthand for*

$$\text{PR}\left(\delta^M(\vec{c}_1, \vec{c}_2) \leq \Theta \mid \delta^M(\vec{r}_1, \vec{r}_2) = x\right).$$

Then,

$$\text{PR}_{\delta^M \leq \Theta}(x) = \frac{\sum_{i=\lceil \frac{x-\Theta}{2} \rceil}^{\lfloor \frac{x+\Theta}{2} \rfloor} \binom{x}{i}}{2^x} \quad (5)$$

where $\Theta, x \in \mathbb{Z}_{0 \dots \lceil \frac{k}{b} \rceil}$ such that $\Theta < x$.

PROOF SKETCH 3. If the Hamming/Manhattan distance between \vec{r}_1 and \vec{r}_2 is x , then it means that these two vectors will differ in exactly x bits, as shown below.

$$\begin{array}{l} \vec{r}_1 : \square\square\square \overbrace{111 \dots 1}^a 0 \dots 000 \square\square\square \\ \vec{r}_2 : \square\square\square 000 \dots 0 \underbrace{1 \dots 111}_{x-a} \square\square\square \end{array}$$

Furthermore, if \vec{r}_1 has $\Delta + a$ bits set to 1, then \vec{r}_2 must have $\Delta + (x - a)$ bits set to 1, where Δ denotes the number of matching 1-bits between \vec{r}_1 and \vec{r}_2 and $a \in \{0, \dots, x\}$. Note that when $b = 1$,

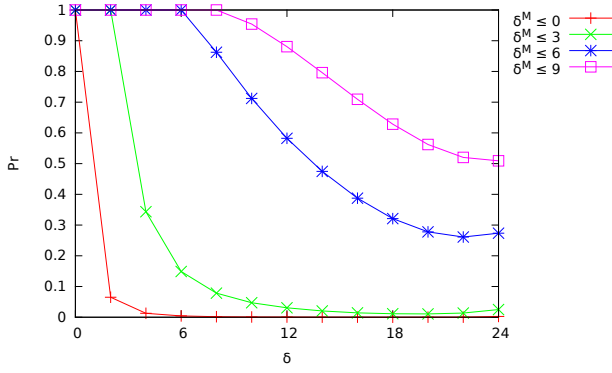


Figure 3: $\text{Pr}_{\delta^M \leq \Theta}$ for $k = 24$ and $b = 6$

the Manhattan distance between \vec{c}_1 and \vec{c}_2 is equal to the difference in the number of 1-bits that \vec{r}_1 and \vec{r}_2 have. Hence,

$$\begin{aligned} \delta^M(\vec{c}_1, \vec{c}_2) &= |\Delta + x - a - (\Delta + a)| \\ &= |x - 2a|. \end{aligned}$$

It is easy to see that there are $(x + 1)$ different configurations:

$$\begin{aligned} a = 0 &\Rightarrow \delta^M(\vec{c}_1, \vec{c}_2) = x \\ a = 1 &\Rightarrow \delta^M(\vec{c}_1, \vec{c}_2) = x - 2 \\ &\vdots \\ a = x - 1 &\Rightarrow \delta^M(\vec{c}_1, \vec{c}_2) = x - 2 \\ a = x &\Rightarrow \delta^M(\vec{c}_1, \vec{c}_2) = x. \end{aligned}$$

Only when $a = \{\lceil \frac{x-\Theta}{2} \rceil, \dots, \lfloor \frac{x+\Theta}{2} \rfloor\}$, will $\delta^M(\vec{c}_1, \vec{c}_2) \leq \Theta$ be satisfied. For each satisfying value of a , the non-matching bits in \vec{r}_1 and \vec{r}_2 can be combined in $\binom{x}{a}$ possible ways. Therefore, there are a total of

$$\sum_{i=\lceil \frac{x-\Theta}{2} \rceil}^{\lfloor \frac{x+\Theta}{2} \rfloor} \binom{x}{i}$$

combinations such that $\delta^M(\vec{c}_1, \vec{c}_2) \leq \Theta$. Since there are 2^x possible combinations in total, the posterior probability in Thm. 3 holds. ■

Using Thm. 3, it is possible to show that when $b = 1$, for all $\Theta < x$ where $\Theta, x \in \mathbb{Z}_{0 \dots (k-2)}$, the following holds:

$$\text{Pr}_{\delta^M \leq \Theta}(x) > \text{Pr}_{\delta^M \leq \Theta}(x+2). \quad (6)$$

Therefore, h_1 is locality-sensitive for $b = 1$. Due to space limitations, we omit the proof of Eqn. 6, but in a nutshell, the proof follows from the fact that going from x to $x+2$, the denominator in Eqn. 5 always increases by a factor of 4, whereas the numerator increases by a factor that is strictly less than 4.

Generalizing Thm. 3 and Eqn. 6 to cases where $b \geq 2$ is more complicated. However, our empirical analyses across multiple values of k and b demonstrate that

$$\text{Pr}_{\delta^M \leq \Theta}(x) \gg \text{Pr}_{\delta^M \leq \Theta}(y)$$

holds when $y \gg x$. For example, Fig. 3 shows $\text{Pr}_{\delta^M \leq \Theta}(x)$ when $k = 24$ and $b = 6$. Fig. 3, along with our empirical evaluations, verify that h_1 is locality sensitive. Thus, combined with a space-filling curve, it can be used to approximate the clustering problem.

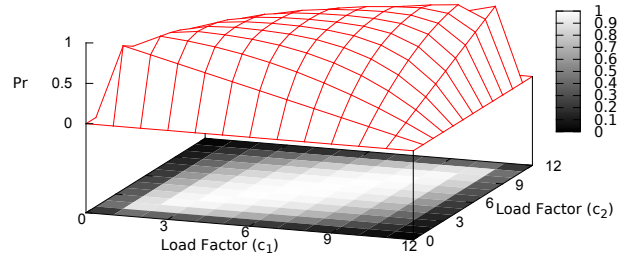


Figure 4: $\text{Pr}(\delta^M \neq \delta)$ for $k = 12, b = 1$ and across varying load factors

5.2 Achieving and Maintaining Tighter Bounds on Tunable-LSH

Next, we demonstrate how it is possible to reduce the approximation error of h_1 . We first define *load factor* of an entry of a record utilization counter.

DEFINITION 2 (LOAD FACTOR). Given a record utilization counter $\vec{c} = (c[0], \dots, c[b-1])$ with size b , the load factor of the i^{th} entry is $c[i]$.

THEOREM 4 (EFFECTS OF GROUPING). Given two record utilization vectors \vec{r}_1 and \vec{r}_2 with size k , let \vec{c}_1 and \vec{c}_2 denote two record utilization counters with size $b = 1$ such that $\vec{c}_1 = h_1(\vec{r}_1)$ and $\vec{c}_2 = h_1(\vec{r}_2)$. Then,

$$\text{Pr} \left(\begin{array}{c} \delta^M(\vec{c}_1, \vec{c}_2) \\ = \delta^M(\vec{r}_1, \vec{r}_2) \end{array} \middle| \begin{array}{c} c_1[0] = l_1 \text{ AND} \\ c_2[0] = l_2 \end{array} \right) = \gamma \quad (7)$$

where

$$\gamma = \frac{\binom{l_{\max}}{l_{\min}} \binom{k}{l_{\max}}}{\binom{k}{l_{\max}} \binom{k}{l_{\min}}} \quad (8)$$

and

$$\begin{aligned} l_{\max} &= \max(l_1, l_2) \\ l_{\min} &= \min(l_1, l_2). \end{aligned}$$

PROOF SKETCH 4. Let \vec{r}_{\max} denote the record utilization vector with the most number of 1-bits among \vec{r}_1 and \vec{r}_2 , and let \vec{r}_{\min} denote the vector with the least number of 1-bits. When $b = 1$, $\delta^M(\vec{c}_1, \vec{c}_2) = \delta(\vec{r}_1, \vec{r}_2)$ holds if and only if the number of 1-bits on which \vec{r}_1 and \vec{r}_2 are aligned is l_{\min} because in that case, both $\delta^M(\vec{c}_1, \vec{c}_2)$ and $\delta(\vec{r}_1, \vec{r}_2)$ are equal to $l_{\max} - l_{\min}$ (note that $\delta^M(\vec{c}_1, \vec{c}_2)$ is always equal to $l_{\max} - l_{\min}$). Assuming that the positions of 1-bits in \vec{r}_{\max} are fixed, there are $\binom{l_{\max}}{l_{\min}}$ possible ways of arranging the 1-bits of \vec{r}_{\min} such that $\delta(\vec{r}_1, \vec{r}_2) = l_{\max} - l_{\min}$. Since the 1-bits of \vec{r}_{\max} can be arranged in $\binom{k}{l_{\max}}$ different ways, there are $\binom{l_{\max}}{l_{\min}} \binom{k}{l_{\max}}$ combinations such that $\delta^M(\vec{c}_1, \vec{c}_2) = \delta(\vec{r}_1, \vec{r}_2)$. Note that in total, the bits of \vec{r}_1 and \vec{r}_2 can be arranged in $\binom{k}{l_{\max}} \binom{k}{l_{\min}}$ possible ways; therefore, Eqns. 7 and 8 describe the posterior probability that $\delta^M(\vec{c}_1, \vec{c}_2) = \delta(\vec{r}_1, \vec{r}_2)$, given $c_1[0] = l_1$ and $c_2[0] = l_2$. ■

According to Eqns. 7 and 8 in Thm. 4, the probability that $\delta^M(\vec{c}_1, \vec{c}_2)$ is an approximation of $\delta(\vec{r}_1, \vec{r}_2)$, but that it is not exactly equal to $\delta(\vec{r}_1, \vec{r}_2)$ is lower for load factors that are close or equal to zero and likewise for load factors that are close or equal to $\lceil \frac{k}{b} \rceil$ (cf., Fig. 4). This property suggests that by carefully choosing f , it is possible to achieve even tighter error bounds for h_1 .

Symbol	Description
begin	natural number between $0 \dots (k - 1)$, initial value is 0
size	natural number between $0 \dots (k - 1)$, keeps track of the number of query access vectors that are currently being maintained, initial value is 0
$H_{k \times ?}$	matrix that contains MIN-HASH values for each query access vector
$S[]$	array of <i>vector</i> (s), one for each MDS query point, that pairs each MDS query point with a <i>random</i> subset of points
$N[]$	array of <i>max-heap</i> (s), one for each MDS query point, that pairs each MDS query point with a set of <i>neighboring</i> points
$X[]$	array of <i>float</i> (s), represents the coordinate (single dimensional) of each MDS query point
$V[]$	array of <i>float</i> (s), represents the current (directional) velocity of each MDS query point

Table 2: Data structures referenced in algorithms

Contrast the matrices in Fig 2b and Fig 2c, which contain the same query access vectors, but the columns are grouped in two different ways²: (i) in Fig. 2b, the grouping is based on the original sequence of execution, and (ii) in Fig. 2c, queries with similar access patterns are grouped together. Fig. 2d and Fig. 2e represent the corresponding record utilization counters for the record utilization vectors in the matrices in Fig. 2b and Fig. 2c, respectively. Take r_3 and r_5 , for instance. Their actual Hamming distance with respect to q_0-q_7 is 8. Now consider the transformed matrices. According to Fig. 2d, the Hamming distance lower bound is 0, whereas according to Fig. 2e, it is 8. Clearly, the bounds in the second representation are closer to the original. The reason is as follows. Even though r_3 and r_5 differ on all the bits for q_0-q_7 , when the bits are grouped as in Fig. 2b, the counts alone cannot distinguish the two bit vectors. In contrast, if the counts are computed based on the grouping in Fig. 2c (which clearly places the 1-bits in separate groups), the counts indicate that the two bit vectors are indeed different.

The observations above are in accordance with Thm. 4. Consequently, we make the following optimization. Instead of randomly choosing a hash function, we construct f such that it maps queries with similar access vectors (i.e., columns in the matrix) to the same hash value. This way, it is possible to obtain record utilization counters with entries that have either very high or very low load factors (cf., Def. 1), thus, decreasing the probability of error (cf., Thm. 4).

We develop a technique to efficiently determine groups of queries with similar access patterns and to adaptively maintain these groups as the access patterns change. Our approach consists of two parts: (i) to approximate the similarity between any two queries, we rely on the MIN-HASH scheme [18], and (ii) to adaptively group similar queries, we develop an incremental version of a multidimensional scaling (MDS) algorithm [48].

MIN-HASH offers a quick and efficient way of approximating the similarity, (more specifically, the Jaccard similarity [41]), between two sets of integers. Therefore, to use it, the query access vectors in our conceptualization need to be translated into a set of positional identifiers that correspond to the records for which the bits in the vector are set to 1.³ For example, according to Fig. 2a, q_1 should be represented with the set $\{0, 5, 6\}$ because r_0 , r_5 and r_6 are the only records for which the bits are set to 1. Note that, we do not need to store the original query access vectors at all. In fact, after the access patterns over a query are determined, we compute and store only its MIN-HASH value. This is important for keeping the memory overhead of our algorithm low.

²Groups are separated by vertical dashed lines.

³In practice, this translation never takes place because the system maintains positional vectors to begin with.

Algorithm 4 Reconfigure-F

Require:

q_t : query access vector produced at time t

Ensure:

Coordinates of MDS points are updated, which are used in determining the outcome of f

```

1: procedure RECONFIGURE-F( $q_t$ )
2:   pos  $\leftarrow$  (begin + size) % k
3:   S[pos].clear()
4:   N[pos].clear()
5:   X[pos]  $\leftarrow$  -0.5 + rand() / RAND-MAX
6:   V[pos]  $\leftarrow$  0
7:   H[pos]  $\leftarrow$  MIN-HASH( $q_t$ )
8:   if size < k then
9:     size += 1
10:  else
11:    begin = (begin + 1) % k
12:  end if
13:  for i  $\leftarrow$  0, i < size, i++ do
14:    x  $\leftarrow$  (begin + i) % k
15:    UPDATE-S-AND-N(x)
16:    UPDATE-VELOCITY(x)
17:  end for
18:  for i  $\leftarrow$  0, i < size, i++ do
19:    x  $\leftarrow$  (begin + i) % k
20:    UPDATE-COORDINATES(x)
21:  end for
22: end procedure

```

Queries with similar access patterns are grouped together using a multidimensional scaling (MDS) algorithm [45] that was originally developed for data visualization, and has recently been used for clustering [16]. Given a set of points and a distance function, MDS assigns coordinates to points such that their original distances are preserved as much as possible. In one efficient implementation [48], each point is initially assigned a random set of coordinates, but these coordinates are adjusted iteratively based on a spring-force analogy. That is, it is assumed that points exert a force on each other that is proportional to the difference between their actual and observed distances, where the latter refers to the distance that is computed from the algorithm-assigned coordinates. These forces are used for computing the current velocity (V in Table 2) and the approximated coordinates of a point (X in Table 2). The intuition is that, after successive iterations, the system will reach equilibrium, at which point, the approximated coordinates can be reported. Since computing all pairwise distances can be prohibitively expensive, the algorithm relies on a combination of sampling ($S[]$ in Table 2) and maintaining for each point, a list of its nearest neighbours ($N[]$ in Table 2)—only these distances are used in computing the net force acting on a point. Then, the nearest neighbours are updated in each iteration by removing the most distant neighbour of a point and replacing it with a new point from the random sample if the distance between the point and the random sample is smaller than the distance between the point and its most distant neighbour.

This algorithm cannot be used directly for our purposes because it is not incremental. Therefore, we propose a revised MDS algorithm that incorporates the following modifications:

1. In our case, each point in the algorithm represents a query access vector. However, since we are not interested in visualizing these points, but rather clustering them, we configure the algorithm to place these points along a single dimension. Then, by dividing the coordinate space into consecutive regions, we are able to determine similar query access vectors.
2. Instead of computing the coordinates of all of the points at once, our version makes incremental adjustments to the coordinates every time reconfiguration is needed.

Algorithm 5 Hash Function f

Require:

t : sequence number of a query access vector

Ensure:

$f(t)$ is computed and returned

```

1: procedure  $F(t)$ 
2:    $pos \leftarrow t \% k$ 
3:    $(lo, hi) \leftarrow \text{GROUP-BOUNDS}(X[pos])$ 
4:    $coid \leftarrow \text{CENTROID}(lo, hi)$ 
5:   return  $\text{HASH}(coid) \% b$ 
6: end procedure

```

The revised algorithm is given in Algorithm 4. First, the algorithm decides which MDS point to assign to the new query access vector \tilde{q}_i (line 2). It clears the array and the heap data structures containing, respectively, (i) the randomly sampled, and (ii) the neighbouring set of points (lines 3–4). Furthermore, it assigns a random coordinate to the point within the interval $[-0.5, 0.5]$ (line 5), and resets its velocity to 0 (line 6). Next, it computes the MIN-HASH value of \tilde{q}_i and stores it in $H[pos]$ (line 7). Then, it makes two passes over all the points in the system (lines 13–21), while first updating their sample and neighbouring lists (line 15), computing the net forces acting on them based on the MIN-HASH distances and updating their velocities (line 16); and then updating their coordinates (line 20).

The procedures used in the last part are implemented in a similar way as the original algorithm [48]; that is, in line 15, the sampled points are updated, in line 16, the velocities assigned to the MDS points are updated, and in line 20, the coordinates of the MDS points are updated based on these updated velocities. However, our implementation of the UPDATE-VELOCITY procedure (line 16) is slightly different than the original. In particular, in updating the velocities, we use a decay function so that the algorithm forgets “old” forces that might have originated from the elements in $S[]$ and $N[]$ that have been assigned to new query access vectors in the meantime. Note that unless one keeps track of the history of all the forces that have acted on every point in the system, there is no other way of “undoing” or “forgetting” these “old” forces.

Given the sequence number of a query access vector (t), the outcome of the hash function f is determined based on the coordinates of the MDS point that had previously been assigned to the query access vector by the RECONFIGURE procedure. To this end, the coordinate space is divided into b groups containing points with consecutive coordinates such that there are at most $\lceil \frac{k}{b} \rceil$ points in each group. Then, one option is to use the group identifier, which is a number in $\mathbb{Z}_0 \dots b-1$, as the outcome of f , but there is a problem with this naïve implementation. Specifically, we observed that even though the *relative* coordinates of MDS points within the “same” group may not change significantly across successive calls to the RECONFIGURE procedure, points within a group, as a whole, may shift. This is an inherent (and in fact, a desirable) property of the incremental algorithm. However, the problem is that there may be far too many cases where the group identifier of a point changes just because the absolute coordinates of the group have changed, even though the point continues to be part of the “same” group. To solve this problem, we rely on a method of computing the centroid within a group by taking the MIN-HASH of the identifiers of points within that group such that these centroids rarely change across successive iterations. Then, we rely on the identifier of the centroid, as opposed to its coordinates, to compute the group number, hence, the outcome of f . The pseudocode of this procedure is given in Algorithm 5.

We make one last observation. Internally, MIN-HASH uses multiple hash functions to approximate the degree to which two sets

	0	1	2	3	4	5
$t = 0$	□	□	□	∅		
$t = \lceil \frac{k}{3} \rceil$		□	□	□	∅	
$t = \lceil \frac{2k}{3} \rceil$			□	□	□	∅
$t = k$				□	□	□
$t = \lceil \frac{4k}{3} \rceil$	□	∅			□	□
$t = \lceil \frac{5k}{3} \rceil$	□	□	∅			□

Figure 5: Assuming $b = 3$, □ indicates the allowed locations at each time tick, and ∅ indicates the counter to be reset.

are similar [18]. It is also known that increasing the number of internal hash functions used (within MIN-HASH) should increase the overall accuracy of the MIN-HASH scheme. However, as un-intuitive as it may seem, in our approach, we use only a single hash function within MIN-HASH, yet, we are still able to achieve sufficiently high accuracy. The reason is as follows. Recall that Algorithm 4 relies on multiple pairwise distances to position every point. Consequently, even though individual pairwise distances may be inaccurate (because we are just using a single hash function within MIN-HASH), collectively the errors are cancelled out, and points can be positioned accurately on the MDS coordinate space.

5.3 Resetting Old Entries in Record Utilization Counters

Once the group identifier is computed (cf., Algorithm 5), it should be straightforward to update the record utilization counters (cf., line 8 in Algorithm 2). However, unless we maintain the original query access vectors, we have no way of knowing which counters to decrement when a query access vector becomes stale, as maintaining these original query access vectors is prohibitively expensive. Therefore, we develop a more efficient scheme in which old values can also be removed from the record utilization counters.

Instead of maintaining b entries in every record utilization counter, we maintain twice as many entries ($2b$). Then, whenever the TUNE procedure is called, instead of directly using the outcome of $f(t)$ to locate the counters to be incremented, we map $f(t)$ to a location within an “allowed” region of consecutive entries in the record utilization counter (cf., line 8 in Algorithm 2). At every $\lceil \frac{k}{b} \rceil^{\text{th}}$ iteration, this allowed region is shifted by one to the right, wrapping back to the beginning if necessary. Consider Fig. 5. Assuming that $b = 3$ and that at time $t = 0$ the allowed region spans entries from 0 to $(b - 1)$, at time $t = \lceil \frac{k}{b} \rceil$, the region will span entries from 1 to b ; at time $t = k$, the region will span entries from b to $2b - 1$; and at time $t = \lceil \frac{4k}{b} \rceil$, the region will span entries 0 and those from $b + 1$ to $2b - 1$.

Since $f(t)$ produces a value between 0 and $b - 1$ (inclusive), whereas the entries are numbered from 0 to $2b - 1$ (inclusive), the RECONFIGURE procedure in Algorithm 2 uses $f(t)$ as follows. If the outcome of $f(t)$ directly corresponds to a location in the allowed region, then it is used. Otherwise, the output is incremented by b (cf., line 8 in Algorithm 2). Whenever the allowed region is shifted to the right, it may land on an already incremented entry. If that is the case, that entry is reset, thereby allowing “old” values forgotten (cf., line 11 in Algorithm 2). These are shown by ∅ in Fig. 5. This scheme guarantees any query access pattern that is less than k steps old is remembered, while any query access pattern that is more than $2k$ old is forgotten.

6. EXPERIMENTAL EVALUATION

	CDB [ICDE'15]	CDB [TUNABLE-LSH]	RDF-3x	VOS [6.1]	VOS [7.1]	MonetDB	4Store
WatDiv 10M	4.7	19.4	18.8	44.0	24.5	17.0	93.0
WatDiv 100M	40.4	42.0	71.4	210.3	96.4	62.7	767.2

Table 3: Query execution time, geometric mean (milliseconds) [11]

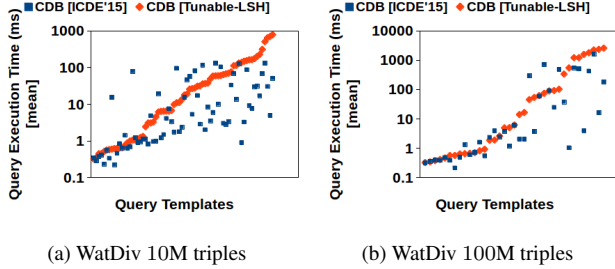


Figure 6: Comparison of chameleon-db implemented using a hierarchical clustering algorithm and with TUNABLE-LSH

In this section, we evaluate TUNABLE-LSH in three sets of experiments. First, we evaluate it within *chameleon-db*, our prototype RDF data management system [10]. Second, we evaluate it within a hashtable implementation since hashtables are used extensively in RDF data management systems. Finally, we evaluate TUNABLE-LSH in isolation, to understand how it behaves under different types of workloads. All experiments are performed on a commodity machine with AMD Phenom II \times 4 955 3.20 GHz processor, 16 GB of main memory and a hard disk drive with 100 GB of free disk space. The operating system is Ubuntu 12.04 LTS.

6.1 Tunable-LSH in chameleon-db

The first experiment evaluates TUNABLE-LSH within *chameleon-db*, which is our prototype RDF data management system [10]. In particular, in earlier work, we had introduced a hierarchical clustering algorithm for grouping RDF triples into what we call *group-by-query clusters* [11]. In this evaluation, we replace that hierarchical clustering algorithm with TUNABLE-LSH, and study its implications on the end-to-end query performance, keeping the same experimental configuration. For completeness, we quote our description of the experimental setup from our previous paper:

“For our evaluations, we [primarily] use the Waterloo SPARQL Diversity Test Suite (WatDiv) because it facilitates the generation of test cases that are far more diverse than any of the existing benchmarks [8]. In this regard, we use the WatDiv *data generator* to create two datasets: one with 10 million RDF triples and another with 100 million RDF triples (we observe that systems under test (SUT) load data into main memory on the smaller dataset whereas at 100M triples, SUTs perform disk I/O). Then, using the WatDiv *query template generator*, we create 125 query templates and instantiate each query template with 100 queries, thus, obtaining 12500 queries.” [11]

We compare our approach with chameleon-db implemented with the hierarchical clustering algorithm (abbreviated CDB [ICDE'15]) and “five popular systems, namely, RDF-3x [51], MonetDB [37], 4Store [33] and Virtuoso Open Source (VOS) versions 6.1 [27] and

7.1 [26]. RDF-3x follows the single-table approach and creates multiple indexes; MonetDB is a column-store, where RDF data are represented using vertical partitioning [1]; and the last three systems are industrial systems. Both 4Store and VOS group and index data primarily based on RDF predicates, but VOS 6.1 is a row-store whereas VOS 7.1 is a column-store. We configure these systems so that they make as much use of the available main memory as possible.” [11]

“We evaluate each system independently on each query template. Specifically, for each query template, we first warm up the system by executing the workload for that query template once (i.e., 100 queries). Then, we execute the same workload five more times (i.e., 500 queries). We report average query execution time over the last five workloads.” [11]

“Our prototype starts with a completely segmented clustering, where each cluster consists of a single triple.” [11] However, “after the execution of the 100th query, we allow the storage advisor to compute a better group-by-query clustering” [11] using either the hierarchical clustering algorithm in [11] or TUNABLE-LSH.

Our experiments indicate that on average, the time to compute the group-by-query clusters has decreased by an order of magnitude with the introduction of TUNABLE-LSH. For example, for the 100M triples dataset, it took 317.6 milliseconds on (geometric) average to compute the group-by-query clusters using the hierarchical clustering algorithm in [11], whereas with TUNABLE-LSH, it takes only about 26.1 milliseconds. This is due to the approximate nature of TUNABLE-LSH. As shown in Table 3 and in Fig. 6, this approximation has a slight impact on query performance, but for the 100M triples dataset, CDB is still significantly faster than the other RDF data management systems. There is one apparent reason for this: TUNABLE-LSH is an approximate method, and therefore, the generated group-by-query clusters are not perfect. To verify this hypothesis, we studied the logs generated during our experiments, which revealed the following: using the group-by-query clustering in [11], chameleon-db’s query engine was able to execute 64.8% of the queries without any decomposition (a property that chameleon-db’s query optimizer is trying to achieve [11]), whereas, the group-by-query clustering computed using TUNABLE-LSH resulted in only 27.1% of the queries to be executed without decomposition. Of course, it is possible to improve chameleon-db’s query optimizer, but that is a topic for future research.

This trade-off between the clustering overhead and the query execution time suggests that for RDF workloads that are too dynamic to be predicted and sampled upfront, it might be desirable to have frequent clustering steps, in which case, using TUNABLE-LSH is a much better option because of its lower overhead.

6.2 Self-Clustering Hashtable

The second experiment evaluates an in-memory hashtable that we developed that uses TUNABLE-LSH to dynamically cluster records in the hashtable. Hashtables are commonly used in RDF data management systems. For example, the dictionary in an RDF data management system, which maps integer identifiers to URIs or literals (and vice versa), is often implemented as a hashtable [1, 26, 58]. Secondary indexes can also be implemented as hashtables, whereby the hashtable acts as a key-value store and maps tuple identifiers to the content of the tuples. In fact, in our own prototype RDF system, *chameleon-db*, all the indexes are secondary (dense) indexes because instead of relying on any sort order inherent in the data, we rely on the notion of group-by-query clusters, in which RDF triples are ordered purely based on the workload [10, 11].

The hashtable interface is very similar to that of a standard hashtable; except that users are given the option to mark the beginning

⁴<http://db.uwaterloo.ca/watdiv/stress-workloads.tar.gz>

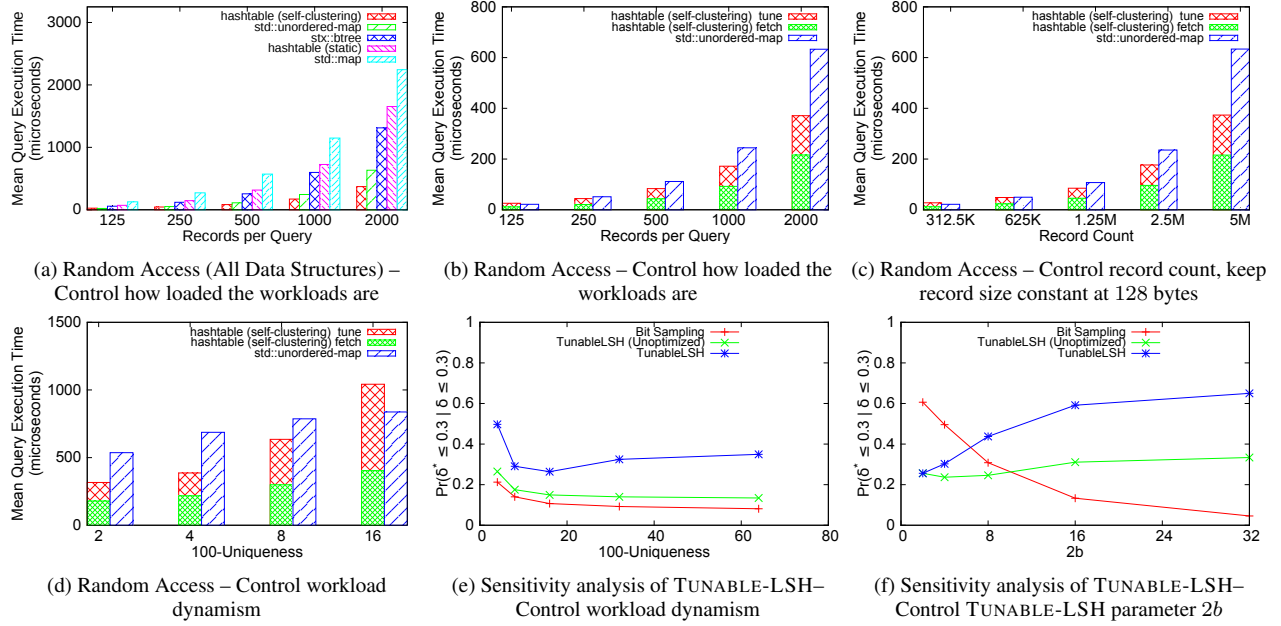


Figure 7: Experimental evaluation of TUNABLE-LSH in a self-clustering hashtable and the sensitivity analysis of TUNABLE-LSH

and end of queries. This information is used to dynamically cluster records such that those that are co-accessed across similar sets of queries also become physically co-located. All of the clustering and re-clustering is transparent to the user, hence, we call this the *self-clustering hashtable*.

The self-clustering hashtable has the following advantages and disadvantages: Compared to a standard hashtable that tries to avoid hash-collisions, it deliberately co-locates records that are accessed together. If the workloads favour a scenario in which many records are frequently accessed together, then we can expect the self-clustering hashtable to have improved fetch times due to better CPU cache utilization, prefetching, etc. [5]. On the other hand, these optimizations come with three types of overhead. First, every time a query is executed, TUNABLE-LSH needs to be updated (cf., Algorithms 2 and 4). Second, compared to a standard hashtable in which the physical address of a record is determined solely using the underlying hash function (which is deterministic throughout the entire workload), in our case, the physical address of a record needs to be maintained dynamically because the underlying hash function is not deterministic (i.e., it is also changing dynamically throughout the workload). Consequently, there is the overhead of going to a lookup table and retrieving the physical address of a record. Third, physically moving records around in the storage system takes time—in fact, this is often an expensive operation. Therefore, the objective of this set of experiments is twofold: (i) to evaluate the circumstances under which the self-clustering hashtable outperforms other popular data structures, and (ii) to understand when the tuning overhead may become a bottleneck. Consequently, we report the end-to-end query execution times, and if necessary, break it down into the time to (i) *fetch* the records, and (ii) *tune* the data structures (which includes all types of overhead listed above).

In our experiments, we compare the self-clustering hashtable to popular implementations of three data structures. Specifically, we use: (i) *std::unordered_map*⁵, which is the C++ standard library

implementation of a hashtable, (ii) *std::map*⁶, which is the C++ standard library implementation of a red-black tree, and (iii) *stx::btree*⁷, which is an open source in-memory B+ tree implementation. As a baseline, we also include a static version of our hashtable, i.e., one that does not rely on TUNABLE-LSH.

We consider two types of workloads: one in which records are accessed *sequentially* and the other in which records are accessed *randomly*. Each workload consists of 3000 queries that are synthetically generated using WatDiv [8]. For each data structure, we measure the end-to-end workload execution time and compute the mean query execution time by dividing the total workload execution time by the number of queries in the workload.

Queries in these workloads consist of changing query access patterns, and in different experiments, we control different parameters such as the number of records that are accessed by queries on average, the rate at which the query access patterns change in the workload, etc. We repeat each experiment 20 times over workloads that are randomly generated with the same characteristics (e.g., average number of records accessed by each query, how fast the workload changes, etc.) and report averages across these 20 runs. We do not report standard errors because they are negligibly small and they do not add significant value to our results.

For the sequential case, *stx::btree* and *std::map* outperform the hashtables, which is expected because once the first few records are fetched from main-memory, the remaining ones can already be prefetched into the CPU cache (due to the predictability of the sequential access pattern). Therefore, for the remaining part, we focus on the random access scenario, which is more common in RDF data management systems, and which can be a bottleneck even in systems like RDF-3x [51] that have clustered indexes over all permutations of attributes. For more examples and a thorough explanation, we refer the reader to [9].

In this experiment, we control the number of records that a query needs to access (on average), where each record is 128 bytes. Fig. 7a

⁵http://www.cplusplus.com/reference/unordered_map/unordered_map/

⁶<http://www.cplusplus.com/reference/map/map/>

⁷<https://panthema.net/2007/stx-btree/>

compares all the data structures with respect to their end-to-end (mean) query execution times. Three observations stand out: First, in the random access case, the self-clustering hashtable as well as the standard hashtable perform much better than the other data structures, which is what would be expected. This observation holds also for the subsequent experiments, therefore, for presentation purposes, we do not include these data structures in Fig. 7b–7d. Second, the baseline static version of our hashtable (i.e., without TUNABLE-LSH) performs much worse than the standard hashtable, even worse than a B+ tree. This suggests that our implementation can be optimized further, which might improve the performance of the self-clustering hashtable as well (this is left as future work). Third, as the number of records that a query needs to access increases, the self-clustering hashtable outperforms all the other data structures, which verifies our initial hypothesis.

For the same experiment above, Fig. 7b focuses on the self-clustering hashtable versus the standard hashtable, and illustrates why the performance improvement is higher (for the self-clustering hashtable) for workloads in which queries access more records. Note that while the *fetch* time of the self-clustering hashtable scales proportionally with respect to `std::unordered_map`, the *tune* overhead is proportionally much lower for workloads in which queries access more records. This is because with increasing “records per query count”, records can be re-located in batches across the pages in main-memory as opposed to moving individual records around.

Next, we keep the average number of records that a query needs to access constant at 2000, but control the number of records in the database. As in the previous experiment, each record is 128 bytes. As illustrated in Fig. 7c, increasing the number of records in the database (i.e., scaling-up) favours the self-clustering hashtable. The reason is that, when there are only a few records in the database, the records are likely clustered to begin with. We repeat the same experiment, but this time, by controlling the record size and keeping the database size constant at 640 megabytes. Surprisingly, the relative improvement with respect to the standard hashtable remains more or less constant, which indicates that the improvement is largely dominated by the size of the database, and increasing it is to the advantage of the self-clustering hashtable.

Finally, we evaluate how sensitive the self-clustering hashtable is to the dynamism in the workloads. Note that for the self-clustering hashtable to be useful at all, the workloads need to be predictable—at least to a certain extent. That is, if records are physically clustered but are never accessed in the future, then all those clustering efforts are wasted. To verify this hypothesis, we control the expected number of query clusters (i.e., queries with similar but not exactly the same access vectors) in any 100 consecutive queries in the workloads that we generate. Let us call this property of the workload, its 100-*Uniqueness*. Fig. 7d illustrates how the *tuning* overhead starts to become a bottleneck as the workloads become more and more dynamic, to the extent of being completely unique, i.e., each query accesses a distinct set of records.

6.3 Sensitivity Analysis of Tunable-LSH

In the final set of experiments, we evaluate the sensitivity of TUNABLE-LSH in isolation, that is, without worrying about how it affects physical clustering, and compare it to three other hash functions: (i) a standard non-locality sensitive hash function⁸, (ii) bit-sampling, which is known to be locality-sensitive for Hamming distances [40], and (iii) TUNABLE-LSH without the optimizations discussed in Section 5. These comparisons are made across workloads with different characteristics (i.e., dense vs. sparse, dynamic

⁸<http://en.cppreference.com/w/cpp/utility/hash>

vs. stable, etc.) where parameters such as the average number of records accessed per query and the expected number of query clusters within any 100-consecutive sequence of queries in the workload are controlled.

Our evaluations indicate that TUNABLE-LSH generally outperforms its alternatives. Due to space considerations, we cannot present all of our results in detail. Therefore, we will summarize our most important observations.

Fig. 7e shows how the probability that *the evaluated hash functions place records with similar utilization vectors to nearby hash values* changes as the workloads become more and more dynamic. In computing these probabilities, both the original distances (i.e., δ) and the distances over the hashed values (i.e., δ^*) are normalized with respect to the maximum distance in each geometry. As illustrated in Fig. 7e, TUNABLE-LSH achieves higher probability even when the workloads are dynamic. The unoptimized version of TUNABLE-LSH behaves more or less like a static locality-sensitive hash function, such as bit sampling, which is an expected result because TUNABLE-LSH cannot achieve high accuracy without the workload-sensitive arrangement introduced in Section 5. It is also important to emphasize that even in that case TUNABLE-LSH is no worse than a standard LSH scheme, which is aligned with the theorems in Section 5.1. We have not included the results on the standard non-locality sensitive hash function, because, as one might guess, it has a probability distribution that is completely unparallelled to our clustering objectives.

Fig. 7f demonstrates how the choice of b (or $2b$ as described in Section 5.3) affects the accuracy of TUNABLE-LSH. Having a higher b implies less and less undesirable collisions of query access vectors, hence, a higher accuracy. On the other hand, for bit sampling, the ideal number of samples is equal to the query clusters in the workload, thus, increasing b , which corresponds to the number of bits that are sampled, might result in oversampling and therefore, lower accuracy. For example, consider two record utilization vectors 1001 and 0001 with Hamming distance 1. If only 1 bit is sampled, there is $\frac{3}{4}$ probability that these two vectors will be hashed to the same value. On the other hand, if 2 bits are sampled, the probability drops to $\frac{1}{2}$.

7. CONCLUSIONS AND FUTURE WORK

In this paper, we introduce TUNABLE-LSH, which is a locality-sensitive hashing scheme, and demonstrate its use in clustering records in an RDF data management system. In particular, we keep track of the fragmented records in the database and use TUNABLE-LSH to decide, in constant-time, where a record needs to be placed in the storage system. TUNABLE-LSH takes into account the most recent query access patterns over the database, and uses this information to auto-tune such that records that are accessed across similar sets of queries are hashed as much as possible to the same or nearby pages in the storage system. This property distinguishes TUNABLE-LSH from existing locality-sensitive hash functions, which are static. Our experiments with (i) a version of our prototype RDF data management system, *chameleon-db*, that uses TUNABLE-LSH, (ii) a hashtable that relies on TUNABLE-LSH to dynamically cluster its records, and (iii) workloads that rigorously test the sensitivity of TUNABLE-LSH verify the potential benefits of TUNABLE-LSH.

As future work, it would be beneficial to answer the following questions. First, the assumption that the last k queries are representative of the future queries in the workload can be relaxed. As outlined in [9], the issue of deciding “when and based on what information to tune the physical design” of our system still remains

an open problem. Second, as our experiments indicate, query optimization in *chameleon-db* has significant room for improvement. We need techniques that can handle more approximate group-by-query clusters such as those generated by TUNABLE-LSH. Third, we believe that TUNABLE-LSH can be used in a more general setting than just RDF systems. In fact, it should be possible to extend the idea of the self-clustering in-memory hashtable that we have implemented to a more general, distributed key-value store.

8. REFERENCES

- [1] D. J. Abadi, A. Marcus, S. R. Madden, and K. Hollenbach. SW-Store: a vertically partitioned DBMS for semantic web data management. *VLDB J.*, 18:385–406, 2009.
- [2] M. Ackerman, S. Ben-David, and D. Loker. Towards property-based classification of clustering paradigms. In *Advances in Neural Information Proc. Systems 23, Proc. 24th Annual Conf. on Neural Information Proc. Systems*, pages 10–18, 2010.
- [3] C. C. Aggarwal. A survey of stream clustering algorithms. In *Data Clustering: Algorithms and Applications*, pages 231–258. 2013.
- [4] S. Agrawal, S. Chaudhuri, and V. R. Narasayya. Automated selection of materialized views and indexes in SQL databases. In *Proc. 26th Int. Conf. on Very Large Data Bases*, pages 496–505, 2000.
- [5] A. Ailamaki, D. J. DeWitt, M. D. Hill, and D. A. Wood. DBMSs on a modern processor: Where does time go? In *Proc. 25th Int. Conf. on Very Large Data Bases*, pages 266–277, 1999.
- [6] R. Al-Harbi, Y. Ebrahim, and P. Kalnis. Phd-store: An adaptive SPARQL engine with dynamic partitioning for distributed RDF repositories. *CoRR*, abs/1405.4979, 2014.
- [7] G. Aluç, D. DeHaan, and I. T. Bowman. Parametric plan caching using density-based clustering. In *Proc. 28th Int. Conf. on Data Engineering*, pages 402–413, 2012.
- [8] G. Aluç, O. Hartig, M. T. Özsu, and K. Daudjee. Diversified stress testing of rdf data management systems. pages 197–212, 2014.
- [9] G. Aluç, M. T. Özsu, and K. Daudjee. Workload matters: Why RDF databases need a new design. *Proc. VLDB Endowment*, 7(10):837–840, 2014.
- [10] G. Aluç, M. T. Özsu, K. Daudjee, and O. Hartig. *chameleon-db: a workload-aware robust RDF data management system*. Technical Report CS-2013-10, University of Waterloo, 2013.
- [11] G. Aluç, M. T. Özsu, K. Daudjee, and O. Hartig. Executing queries over schemaless RDF databases. In *Proc. 31st Int. Conf. on Data Engineering*, 2015. Forthcoming.
- [12] A. Andoni and P. Indyk. Near-optimal hashing algorithms for approximate nearest neighbor in high dimensions. In *Proc. 47th Annual Symp. on Foundations of Computer Science*, pages 459–468, 2006.
- [13] M. Arias, J. D. Fernández, M. A. Martínez-Prieto, and P. de la Fuente. An empirical study of real-world SPARQL queries. *CoRR*, abs/1103.5043, 2011.
- [14] V. Athitsos, M. Potamias, P. Papapetrou, and G. Kollios. Nearest neighbor retrieval using distance-based hashing. In *Proc. 24th Int. Conf. on Data Engineering*, pages 327–336, 2008.
- [15] R. G. Bello, K. Dias, A. Downing, J. J. Feenan, Jr., J. L. Finnerty, W. D. Norcott, H. Sun, A. Witkowski, and M. Ziauddin. Materialized views in oracle. In *Proc. 24th Int. Conf. on Very Large Data Bases*, pages 659–664, 1998.
- [16] B. Bislimovska, G. Aluç, M. T. Özsu, and P. Fraternali. Graph search of software models using multidimensional scaling. In *Proc. of the Workshops of the EDBT/ICDT 2015 Joint Conference*, pages 163–170, 2015.
- [17] M. A. Bornea, J. Dolby, A. Kementsietsidis, K. Srinivas, P. Dantressangle, O. Udrea, and B. Bhattacharjee. Building an efficient RDF store over a relational database. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 121–132, 2013.
- [18] A. Z. Broder. On the resemblance and containment of documents. In *Proc. Compression and Complexity of Sequences*, pages 21–29, 1997.
- [19] A. Z. Broder, M. Charikar, A. M. Frieze, and M. Mitzenmacher. Min-wise independent permutations. *J. Comput. Syst. Sci.*, 60(3):630–659, 2000.
- [20] N. Bruno and S. Chaudhuri. To tune or not to tune? a lightweight physical design alerter. In *Proc. 32nd Int. Conf. on Very Large Data Bases*, pages 499–510, 2006.
- [21] J. J. Carroll, I. Dickinson, C. Dollin, D. Reynolds, A. Seaborne, and K. Wilkinson. Jena: implementing the semantic web recommendations. In *Proc. 13th Int. World Wide Web Conf. - Alternate Track Papers & Posters*, pages 74–83, 2004.
- [22] S. Ceri, S. B. Navathe, and G. Wiederhold. Distribution design of logical database schemas. *IEEE Trans. Softw. Eng.*, 9(4):487–504, 1983.
- [23] M. Charikar. Similarity estimation techniques from rounding algorithms. In *Proc. 34th Annual ACM Symp. on Theory of Computing*, pages 380–388, 2002.
- [24] S. Chaudhuri and V. Narasayya. Self-tuning database systems: a decade of progress. In *Proc. 33rd Int. Conf. on Very Large Data Bases*, pages 3–14, 2007.
- [25] M. Datar, N. Immorlica, P. Indyk, and V. S. Mirrokni. Locality-sensitive hashing scheme based on p-stable distributions. In *Proc. 20th Annual Symp. on Computational Geometry*, pages 253–262, 2004.
- [26] O. Erling. Virtuoso, a hybrid RDBMS/graph column store. *IEEE Data Eng. Bull.*, 35(1):3–8, 2012.
- [27] O. Erling and I. Mikhailov. RDF support in the Virtuoso DBMS. *Networked Knowledge-Networked Media*, pages 7–24, 2009.
- [28] H. Ferhatosmanoglu, E. Tuncel, D. Agrawal, and A. El Abbadi. Approximate nearest neighbor searching in multimedia databases. In *Proc. 17th Int. Conf. on Data Engineering*, pages 503–511, 2001.
- [29] L. Galarraga, K. Hose, and R. Schenkel. Partout: a distributed engine for efficient RDF processing. In *Proc. 23rd Int. World Wide Web Conf., Companion Volume*, pages 267–268, 2014.
- [30] A. Gionis, P. Indyk, and R. Motwani. Similarity search in high dimensions via hashing. In *Proc. 25th Int. Conf. on Very Large Data Bases*, pages 518–529, 1999.
- [31] F. Goasdoué, K. Karanasos, J. Leblay, and I. Manolescu. View selection in semantic web databases. *Proc. VLDB Endowment*, 5(2):97–108, 2011.
- [32] R. W. Hamming, editor. *Coding and Information Theory*. Prentice-Hall, Englewood Cliffs, 1986.
- [33] S. Harris, N. Lamb, and N. Shadbolt. 4store: The design and implementation of a clustered RDF store. In *Proc. 5th Int.*

Workshop on Scalable Semantic Web Knowledge Base Systems, pages 81–96, 2009.

- [34] A. Harth, J. Umbrich, A. Hogan, and S. Decker. Yars2: A federated repository for querying graph structured data from the web. In *Proc. 6th Int. Semantic Web Conference*, pages 211–224, 2007.
- [35] K. Hose and R. Schenkel. WARP: workload-aware replication and partitioning for RDF. In *Proc. Workshops of the 29th IEEE Int. Conf. on Data Engineering*, pages 1–6, 2013.
- [36] M. E. Houle and J. Sakuma. Fast approximate similarity search in extremely high-dimensional data sets. In *Proc. 21st Int. Conf. on Data Engineering*, pages 619–630, 2005.
- [37] S. Idreos, F. Groffen, N. Nes, S. Manegold, K. S. Mullender, and M. L. Kersten. MonetDB: Two decades of research in column-oriented database architectures. *IEEE Data Eng. Bull.*, 35(1):40–45, 2012.
- [38] S. Idreos, M. L. Kersten, and S. Manegold. Database cracking. In *Proc. 3rd Biennial Conf. on Innovative Data Systems Research*, pages 68–78, 2007.
- [39] S. Idreos, S. Manegold, H. A. Kuno, and G. Graefe. Merging what’s cracked, cracking what’s merged: Adaptive indexing in main-memory column-stores. *PVLDB*, 4(9):585–597, 2011.
- [40] P. Indyk and R. Motwani. Approximate nearest neighbors: Towards removing the curse of dimensionality. In *Proc. 30th Annual ACM Symp. on Theory of Computing*, pages 604–613, 1998.
- [41] P. Jaccard. The distribution of flora in the alpine zone. *New phytologist*, 11(2):37–50, 1912.
- [42] A. K. Jain, M. N. Murty, and P. J. Flynn. Data clustering: a review. 31:264–323, 1999.
- [43] M. Kirchberg, R. K. L. Ko, and B.-S. Lee. From linked data to relevant data – time is the essence. *CoRR*, abs/1103.5046, 2011.
- [44] E. F. Krause, editor. *Taxicab Geometry: An Adventure in Non-Euclidean Geometry*. Dover, New York, 1986.
- [45] J. B. Kruskal. Multidimensional scaling by optimizing goodness of fit to a nonmetric hypothesis. *Psychometrika*, 29:1–27, 1964.
- [46] V. I. Levenshtein. *Soviet Physics Doklady*, 10(8):707–710, 1966.
- [47] S. Lightstone, T. J. Teorey, and T. P. Nadeau. *Physical Database Design: the database professional’s guide to exploiting indexes, views, storage, and more*. Morgan Kaufmann, 2007.
- [48] A. Morrison, G. Ross, and M. Chalmers. Fast multidimensional scaling through sampling, springs and interpolation. *Information Visualization*, 2(1):68–77, 2003.
- [49] G. M. Morton. A computer oriented geodetic data base; and a new technique in file sequencing. Technical report, IBM Ltd., Ottawa, Canada, 1966.
- [50] F. F.-H. Nah. A study on tolerable waiting time: how long are Web users willing to wait? *Behaviour & IT*, 23(3):153–163, 2004.
- [51] T. Neumann and G. Weikum. The RDF-3X engine for scalable management of RDF data. *VLDB J.*, 19(1):91–113, 2010.
- [52] T. Neumann and G. Weikum. x-RDF-3X: fast querying, high update rates, and consistency for RDF databases. *Proc. VLDB Endowment*, 3(1):256–263, 2010.
- [53] N. Papailiou, I. Konstantinou, D. Tsoumakos, and N. Koziris. H2RDF: adaptive query processing on RDF data in the cloud. pages 397–400, 2012.
- [54] N. Papailiou, D. Tsoumakos, P. Karras, and N. Koziris. Graph-aware, workload-adaptive SPARQL query caching. 2015. Forthcoming.
- [55] L. Sidirourgos, R. Goncalves, M. Kersten, N. Nes, and S. Manegold. Column-store support for RDF data management: not all swans are white. *Proc. VLDB Endowment*, 1(2):1553–1563, 2008.
- [56] Y. Tao, K. Yi, C. Sheng, and P. Kalnis. Efficient and accurate nearest neighbor and closest pair search in high-dimensional space. *ACM Trans. Database Syst.*, 35(3), 2010.
- [57] C. Weiss, P. Karras, and A. Bernstein. Hexastore: sextuple indexing for semantic web data management. *Proc. VLDB Endowment*, 1(1):1008–1019, 2008.
- [58] K. Wilkinson. Jena property table implementation. Technical Report HPL-2006-140, HP-Labs, 2006.
- [59] D. C. Zilio, J. Rao, S. Lightstone, G. M. Lohman, A. J. Storm, C. Garcia-Arellano, and S. Fadden. DB2 design advisor: Integrated automatic physical database design. In *Proc. 30th Int. Conf. on Very Large Data Bases*, pages 1087–1097, 2004.
- [60] L. Zou, J. Mo, D. Zhao, L. Chen, and M. T. Özsu. gStore: Answering SPARQL queries via subgraph matching. *Proc. VLDB Endowment*, 4(1):482–493, 2011.